arXiv:1603.03888v1 [cs.DC] 12 Mar 2016

Performance Evaluation of Unified Parallel C for Molecular Dynamics

Kamran Idrees[1], Christoph Niethammer[1], Aniello Esposito[2]
and Colin W. Glass[1]

[1] High Performance Computing Center Stuttgart (HLRS), Stuttgart, Germany
idrees@hlrs.de, niethammer@hlrs.de, glass@hlrs.de
[2] Cray Inc., Seattle, WA, U.S.A.
esposito@cray.com

Abstract

Partitioned Global Address Space (PGAS) integrates the concepts of shared memory programming and the control of data distribution and locality provided by message passing into a single parallel programming model. The purpose of allying distributed data with shared memory is to cultivate a locality-aware shared memory paradigm. PGAS is comprised of a single shared address space, which is partitioned among threads. Each thread has a portion of the shared address space in local memory and therefore it can exploit data locality by mainly doing computation on local data.
Unified Parallel C (UPC) is a parallel extension of ISO C and an implementation of the PGAS model. In this paper, we evaluate the performance of UPC based on a real-world scenario from Molecular Dynamics.

1 Introduction

Partitioned Global Address Space (PGAS) is a locality-aware distributed shared memory model for Single Program Multiple Data (SPMD) streams. PGAS unites the concept of shared memory programming and distributed data. It provides an abstraction of Global Shared Address Space, where each thread can access any memory location using a shared memory paradigm. The Global Shared Address Space is formed by integrating the portions of the memories on different nodes and the low level communication involved for accessing remote data is hidden from the user. Unified Parallel C (UPC) is an implementation of the PGAS model. The low-level communication in UPC is implemented using light-weight Global-Address Space Networking (GASNet). UPC benefits from the brisk one-sided communication provided by GASNet and thus has a performance advantage over message passing [5].

Molecular Dynamics simulates the interactions between molecules [1]. After the system is initialized, the forces acting on all molecules in the system are calculated. Newton's equations of motion are integrated to advance the positions and velocities of the molecules. The simulation is advanced until the computation of the time evolution of the system is completed for a specified length of time.

In this paper, we evaluate the intra- and inter-node performance of UPC based on a real-world application from Molecular Dynamics, compare it intra-node with OpenMP and show the necessity for manual optimizations by the programmer in order to achieve good performance.

2 Unified Parallel C

Unified Parallel C (UPC) is a parallel extension of ISO C. It is a distributed shared memory programming model that runs in a SPMD fashion, where all threads execute the main program or function. Using UPC constructs, each thread can follow a distinct execution path to work on

different data. UPC threads run independently of each other, the only implied synchronization is at the beginning and at the end of the main function [4]. It is the responsibility of the programmer to introduce necessary synchronization when shared data is accessed by more than one thread. Apart from a global shared address space, UPC also provides private address space for each thread. The private address space is only accessible by the thread whichs owns it. This allows a programmer to intelligently allocate the data in private and shared address spaces. Data which remains local to a thread should be allocated on the private address space. Whereas data which needs to be accessed by multiple UPC threads, should be allocated on the portion of the shared address space of the thread doing most computation on it [4].

UPC accommodates several constructs which allow to allocate data and the thread with affinity to it on the same physical node. UPC also provides constructs to check the locality of data. The programmer needs to identify data as local in order to access it with a local pointer.

UPC utilizes a source to source compiler. The source to source compiler translates UPC code to ANSI C code (with additional code for communication to access remote memory, which is hidden from the user) and links to the UPC run-time system. The UPC run-time system can examine the shared data accesses and perform communication optimizations [5] [4].

3 Molecular Dynamics Code

We ported our in-house Molecular Dynamics code CMD, developed for basic research into high performance computing. CMD features multiple MD data structures, algorithms and parallelization strategies and thus allows for quantitative comparisons between them. Two widely used data structures are implemented - with corresponding algorithms - for the computation of interactions between molecules in the system, "BasicN2" and "MoleculeBlocks". The MoleculeBlocks code has been ported to UPC.

MoleculeBlocks implements a linked cell approach, where the domain is spatially decomposed into cells (of the size cut-off radius) and then the molecules are distributed among these cells. In this algorithm, the distances between the molecules are computed only intra-cell and for neighboring cells. Furthermore, Newton's 3rd law of motion is used to reduce the compute effort by half. Figure 1 shows an example of the MoleculeBlocks algorithm for a 2D domain space. When the interaction between a pair of molecules is computed, the resulting force is written to both molecules. Thus, the centered cell (dark gray), as shown in figure 1, modifies the forces of its own molecules and molecules of its right and lower neighbor cells (gray). Although the use of Newtons 3rd law lessens the computational effort, it raises the requirements regarding synchronization in order to avoid race conditions.

4 Porting Molecular Dynamics Code to UPC

For MoleculeBlocks, the system of molecules (called Phasespace) is spatially decomposed into cells where each cell contains a number of molecules (as shown in figure 2). The cells are then distributed among the UPC threads in a spatially coherent manner (as opposed to the default round-robin fashion) to reduce the communication overhead between the UPC threads. The CMD simulation code is comprised of two parts, (i) phasespace initialization & grid generator and (ii) main simulation loop.

4.1 Phasespace Initialization & Grid Generator

Phasespace initialization involves allocation of the memory dynamically on the global shared space for molecules and cells. A routine for performing the transformation (or mapping) of a spatially coherent cell indices (i,j,k) to consecutive integers (cell IDs) is introduced in the

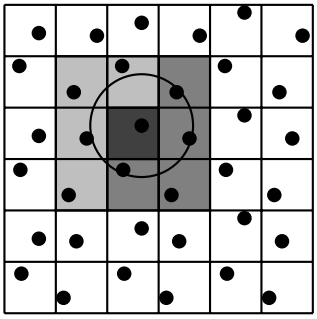

Figure 1: Calculation of interaction between molecules using the `MoleculeBlocks` algorithm.

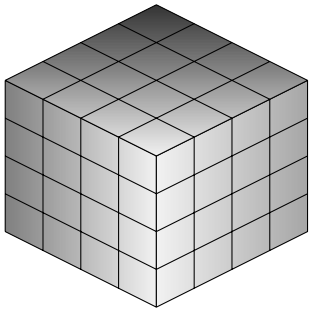

Figure 2: Domain is spatially decomposed into cells and distributed among threads in a spatially coherent manner.

code, which allows UPC shared array to distibute spatially coherent cells among UPC threads by blocking consecutive cell IDs. The Grid generation routine initializes the positions and velocities of the molecules and adds them to the phasespace.

4.2 Main Simulation Loop

All the routines inside the main simulation loop have been implemented in UPC as locality-aware algorithms. Each thread only calculates the different parameters for the cells which reside in its portion of shared space. Due to manual optimization, each thread accesses its cells using local pointer instead of pointer-to-shared.

The Lennard-Jones Force Calculation routine is the most compute intensive routine of the simulation. Here, each thread computes the inter-molecular interactions, for the cells residing in its portion of shared space and with the neighbor cells which may reside locally or remotely. The synchronization among threads is provided through a locking strategy. We have tested different locking granularities, which are "lock per molecule" and "lock per cell". Furthermore, we have also implemented pre-fetching and copy-at-once strategies to reduce the communication between UPC threads. This has a major effect when the UPC code is scaled beyond a single node. With pre-fetching and copy-at-once strategies, a UPC thread pre-fetches the values (of positions and forces) of the molecules of its neighbor cell if the neighbor cell is remote (i.e. does not reside in its local portion of the shared space) to its private space. The function of pre-fetching is implemented using the upc_memget routine. Once a UPC thread has pre-fetched the data, it computes all interactions between the molecules of the two cells and then copies all the calculated forces to the neighbor cell in one go using the upc_memput routine.

In order to calculate the global value of parameters (e.g potential energy), the coordination among threads is done using the reduction function upc_all_reduceT available in the collective library of UPC.

5 Hardware Platforms

The presented benchmarks have been produced on a Cray XE6 system and a Cray XC30 system. The code was built by means of the Berkeley UPC compiler (version 2.16.2) on both systems. The cray compiler had some performance issues which are under investigation. The nodes of the XE6 system feature two AMD Interlagos processors (AMD Opteron(TM) Processor 6272 clocked at 2.10GHz) with 16 integer cores each. Two integer cores share a floating

point unit. On the other hand, the compute nodes of the XC30 system contain two Intel Sandybridge processors (Intel(R) Xeon(R) CPU E5-2670 0 clocked at 2.60GHz) with 8 cores and 8 hyperthreads each. For the present benchmarks no hyperthreads have been used.

For a full documentation and technical specifications of the hardware platforms, the reader is referred to the online material[1].

6 Evaluation

The UPC implementation of CMD is evaluated for the cut-off radius of 3 with the following strategies (a lock always locks the entire cell).

1. *Lock per molecule (LPM)* - Acquire lock for each molecule-molecule interaction
2. *Lock per cell (LPC)* - Acquire lock for each cell-cell interaction
3. *Lock per cell plus prefetching (LPC+)* - Same as lock per cell but with pre-fetching and copy-at-once strategies

The cut-off radius determines the maximum distance for evaluating molecule-molecule interactions. Increasing the cut-off will result in more interactions per molecule and therefore more computational effort. The cell size is equal to the cut-off radius.

In the first UPC implementation of CMD, we did not perform the manual pointer optimizations. The shared address space was always accessed using pointer-to-shared irrespective of the fact whether the data accessed by a UPC thread is local or remote. The test cases of 500, 3,000 and 27,000 molecules without pointer optimizations executed 10 times slower than the current version which incorporates manual pointer optimizations.

The rest of the evaluation is based on the version with manual optimization, a cut-off radius of 3 and with 6,000,000 molecules. Thus, all scaling benchmarks shown here are based on strong scaling. The evaluation metrics are explained in the following subsection.

6.1 Evaluation Metrics

The UPC implementation of CMD is evaluated for all three strategies described above, on the basis of intra- and inter- node performance. For each case, we have taken the average execution time for five program runs.

Intra-Node Performance The intra-node performance compares the UPC and OpenMP implementations of CMD on a single node.

Inter-Node Performance The inter-node performance shows the scaling on multiple nodes. Under populated nodes are used for inter-node results when the number of UPC threads are less than the total available CPU count of 4 nodes. The threads are always equally distributed among the nodes.

6.2 Results

Here we show the execution time of CMD both intra- and inter-node, and compare the execution time with varying locking granularities.

Intra-Node Performance Figure 4 shows the intra-node benchmark for the execution time of UPC implementation of CMD. Clearly, the lock per cell strategy is superior to the lock per molecule. Pre-fetching and copy-at-once has no significant impact on intra-node performance. Figure 3 compares intra-node performance achieved with OpenMP and UPC. UPC performs similarly to OpenMP on a single node. This is a satisfactory result, as the aim is not to provide a better shared memory parallelization, but to use a shared memory paradigm for distributed memory. Having a comparable performance as OpenMP is a good basis.

[1] www.cray.com/Products/Computing/XE/Resources.aspx, www.cray.com/Products/Computing/XC/Resources.aspx

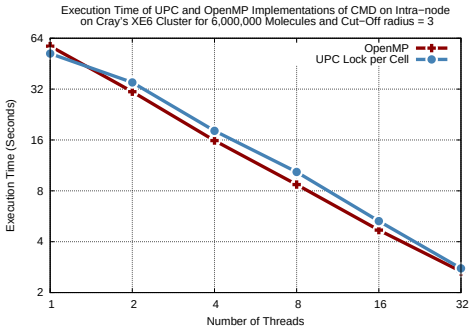

Figure 3: Comparison of intra-node execution time of UPC and OpenMP implementations of CMD on a Cray XE6.

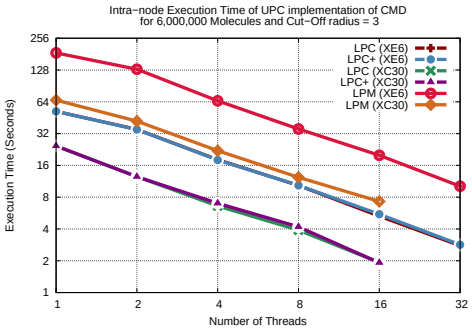

Figure 4: Intra-node execution time of UPC implementation of CMD with different synchronization strategies.

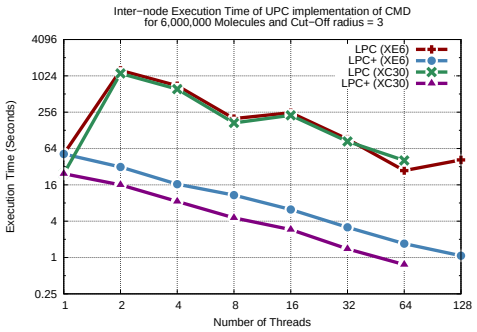

Figure 5: Inter-Node Execution Time of the UPC implementation with different locking strategies and cut-off radius of 3.

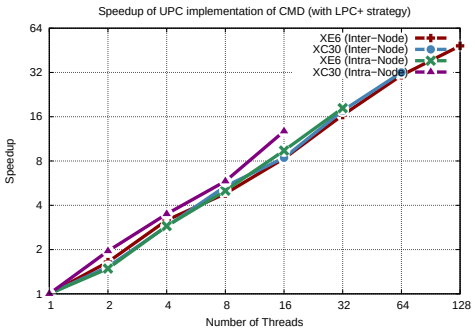

Figure 6: Speedup of UPC implementation of CMD with Lock per cell plus prefetching strategy on intra- and inter-node.

Inter-Node Performance Figure 5 shows the inter-node performance achieved with UPC, using LPC and LPC+. The LPM strategy is disregarded due to its inferior intra-node performance. As can be seen, the LPC strategy shows very poor inter-node performance (2 or more threads). The execution time jumps by a factor of 20+, as soon as inter-node communication comes into play. However, the LPC+ strategy shows a solid scaling behaviour, making this implementation competitive even for production runs. Figure 6 shows the speedup of CMD with LPC+ strategy on intra- and inter-node, for both XE6 and XC30 clusters.

7 Conclusion

As we have shown in this paper, it is possible to implement a competitive distributed memory parallelization using UPC. However, the implementation is far from trivial. There are many pitfalls, leading to significant performance degradations and they are not always obvious.

The first and most elusive problem is the use of pointer-to-shared for local data accesses. As a programmer, one would expect UPC compilers or the run-time to automatically detect local data accesses and perform the necessary optimization. However, this is not the case and the performance degradation for our use case was a staggering factor 10. This suggests that with

the currently available compilers, manual pointer optimization (using local C pointers when a thread has affinity to the data) is mandatory.

The second issue is not discussed in detail here. In a nutshell: the default round robin distribution of shared array elements leads to significant communication traffic in this scenario. Manual optimization was necessary, essentially replacing round robin with a spatially coherent distribution. Thus, the programmer needs to keep the underlying distributed memory architecture in mind.

The third and maybe most disturbing problem is related to communication granularity. The LPM and LPC strategies represent the way one traditionally would approach a shared memory parallelization. As the data is available in shared memory, there is no need to pre-fetch it or to package communication. However, these approaches fail completely when utilized for distributed memory, as can be seen in figure 5.

The good news is, all the above problems can be solved, the bad news is: it requires the programmer to think in terms of distributed memory parallelization. However, this is not the driving idea behind PGAS.

In our view, in order for PGAS approaches to prosper in the future, these issues have to be addressed. Only if better data locality, data distribution and communication pooling is provided automatically by the compiler or run-time will programmers start seeing true benefit. The required information for such improved automatic behaviour is available to the compiler, as the parallelization is achieved purely through UPC data structures and routines.

Acknowledgments This work was supported by the project *HA* which is funded by the German Research Foundation (DFG) under the priority programme "Software for Exascale Computing - SPPEXA" (2013-2015) and the EU project *APOS* which was funded as part of the European Commissions Framework 7.

References

[1] Michael P Allen. Introduction to molecular dynamics simulation. *Computational Soft Matter: From Synthetic Polymers to Proteins*, 23:1–28, 2004.

[2] Cristian Coarfa, Yuri Dotsenko, John Mellor-Crummey, François Cantonnet, Tarek El-Ghazawi, Ashrujit Mohanti, Yiyi Yao, and Daniel Chavarría-Miranda. An evaluation of global address space languages: co-array fortran and unified parallel c. In *Proceedings of the tenth ACM SIGPLAN symposium on Principles and practice of parallel programming*, pages 36–47. ACM, 2005.

[3] Leonardo Dagum and Ramesh Menon. Openmp: an industry standard api for shared-memory programming. *Computational Science & Engineering, IEEE*, 5(1):46–55, 1998.

[4] Tarek El-Ghazawi, William Carlson, Thomas Sterling, and Katherine Yelick. *UPC: distributed shared memory programming*, volume 40. Wiley-Interscience, 2005.

[5] Katherine Yelick, Dan Bonachea, Wei-Yu Chen, Phillip Colella, Kaushik Datta, Jason Duell, Susan L. Graham, Paul Hargrove, Paul Hilfinger, Parry Husbands, Costin Iancu, Amir Kamil, Rajesh Nishtala, Jimmy Su, Michael Welcome, and Tong Wen. Productivity and performance using partitioned global address space languages. In *Proceedings of the 2007 international workshop on Parallel symbolic computation*, PASCO '07, pages 24–32, New York, NY, USA, 2007. ACM